J. Social Biol. Struct. (1985) 8, 13–26

Dynamic models for thought

Ralph H. Abraham

Division of Natural Sciences, University of California, Santa Cruz, CA 95064, USA

This paper is the twelfth of a series, *On Morphodynamics*. Here, we introduce a complex dynamical model for the brain, and present some trial mechanisms for the abstraction and application of ideas. These, based on the concept of the holonomy of a bifurcation diagram, are intended just to indicate the range of possibilities, not as definitive models.

Dedicated to: Aharon Katzir-Katchlasky (1914–1972)

Levels of abstraction

In 'The function of mathematics in the evolution of the noosphere' (Abraham, 1981–OM7†), we proposed a model for a noosphere, an aggregation of conscious organisms. Here we will retreat one step, and consider a single mind. But as this organism is viewed as an aggregation of organs, the model described here is very similar to those proposed previously.

In our geometric model it is convenient to discretize one preferred dimension–*abstraction*. This strategy, explicit in classical Sanskrit (Upanishad) philosophy, replaces a model of dimension n by a finite stack of parallel models, say k planes of dimension $(n - 1)$. If $n = 3$, this is like a deck of cards. A more extreme version of this strategy–division of the geometrical model into a finite set of cells–reduces in one step from dimension n to dimension zero. The cellular structure of biological organisms is an important example of the extreme strategy. We will return to this example later. Now, we consider a stack of k parallel planes of dimension $(n - 1)$ as a model for a single conscious mind. Although the dimension must be large for a reasonable model, we will set $n = 3$ here for the sake of visualization. Likewise, the number of planes, to approximate a continuous scale of abstraction, should be large. The Sanskrit philosophers frequently take $k = 7$. Following Plato, we will take $k = 4$. According to Shear (1977), this structure of consciousness coincides with the development stages of Piaget for the growth of consciousness in children. The four levels of Plato's hierarchy are:

G. The Good–universal archetypes;
M. Mathematics–abstract mental images of archetypes;

†OM7 refers to the seventh article in the series *On Morphodynamics*. Other articles in the series will be referred to in the same manner.

0140-1750/85/010013 + 14 $03.00/0

S. Science/theories/models–mental representation of sensory experiences;
P. Phenomena–data of sensory experience of the phenomenal universe.

We would like to imbed this hierarchical structure in a dynamical model for the mind, so we choose to use *complex dynamical systems* as the basic unit in constructing the model [see Abraham (1983*b*–OM10; or 1983*c*–OM11) for the definitions]. We have proposed a model of this type for the mammalian brain, in 'Vibrations, the realization of form' (Abraham, 1976–OM4). Eventually, our goal in this paper is to study the interaction between two adjacent levels. The model for each level will be based upon a familiar mathematical object; a *simple dynamical scheme,* or in other words, a *dynamical system with controls* (albeit, with a very large number of dimensions). Our theory of interaction would equally well apply to the serially coupled adjacent levels of any complex dynamical system. But for this exposition, we use Mathematics and Science, levels M and S of Plato's model of the mind.

In summary: We construct a partial model for mind with two levels of abstraction:

Level M: Mathematics.
Level S: Science.

In this model, we study interactions based on two-way communications between levels.

Interactions between levels

The movement of information between these two adjacent levels of the model are *abstraction* and *application*. We consider these, one at a time:

Abstraction: information moves from level S to level M in the emergence of an abstract mental image in M, based upon the association, or *aggregation*, of several special cases existing in consciousness on level S.

Application: information moves from level M to level S, through *diffusion* of an abstract image into a region of experiential (or experimental) data to which it can *apply*, or *associate*. as shown in Fig. 1.

We may describe these two interactions, through metaphors, without an explicit scheme for the *representation* of an idea in a neurophysiological model for the mammalian brain. In this section, we describe them in mechanical and informatic metaphors. In a later section we will interpret these metaphors in a specific neurophysiological model.

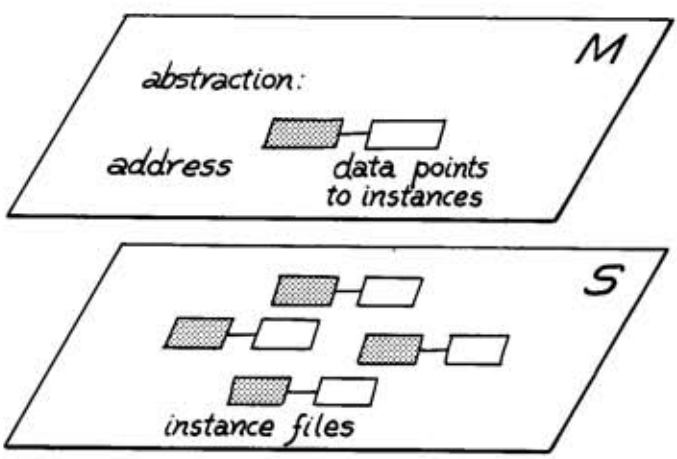

Fig. 1. Informatic metaphor for application.

First, we consider *abstraction*. In the *informatic metaphor*, each instance of a similar theory or model in S may be considered as a *file*, or movable package of information, with a *name*. Actually, the name is a *pointer*, a program which finds the file. The *aggregation* of these instances into a single concept, or file, on a higher level of abstration, only requires moving the files into a eommon *directory*. The directory is just another file, which contains the names of (instructions for finding) the instance files. The structure of this system is a *tree* of information. The *abstraction* operation is the creation of a new directory of level M, containing pointers to all the instance files on level S, which still live there. They have not moved.

Although this metaphor for abstraction is reasonable for a hard-wired computing machine, it is obviously inappropriate for a dynamical network like the mammalian brain. We now transform it into a *mechanical metaphor*. We customarily use this metaphor even when thinking about computer systems, because it seems more natural.

In the mechanical metaphor, we think of each instance (model on level S) exhibiting a given concept (model on level M) as a file folder lying on the desk, rather than as software, attached to an immovable physical address. We associate the files into a directory by moving them physically, into a stack. This stack is not the directory, for it still occupies level S. But as the instances are physically associated to a common address, the coordinates of the pile of file folders, we may use this common address on level S as a name for the abstraction. Thus, the directory is a filename (or pointer program) on level M, which points to a physical address on level S, at which all the instances may be found. This is like the subject catalog in a library.

In the informatic metaphor, the process of forming an abstraction consists of making a list on level M of instance addresses on level S. In the mechanical metaphor, the process consists of moving the instances around on level S, to a common location, and noting its address on level M.

Next, we will describe the *application* process, in both metaphors. In the informatic metaphor, we suppose we have an abstraction on hand, as a directory on level M, containing the names of instances on level S. These instances are files of information, and the names are programs which find these files. The *application process*, for an old application, requires just recognizing the name of the instance, and running the program to link the abstraction to the instance file. But suppose we want to create a *new application*. We must *recognize* (or guess) an existing abstraction which the instance exhibits, and then *name* the instance within the abstraction directory. That is, we must create a program linking the abstraction to the new instance. In the mechancial metaphor, we must recognize the abstraction (the address on level M of a pile of file folders on level S), and then move the new instance onto the pile.

The next few sections will elaborate the informatic metaphor for both processes into a non-local neurophysiological model, and elaborate the mechanical model into a local geometric model. The neurophysiological and the geometric models may be directly related, without the metaphors described above, which are for explanatory purposes only.

The field scheme for organs

A precise model for a mammalian brain, in the context of complex dynamical systems theory or any such discipline, is beyond us at present. What we do have now is an emerging scheme. We have written of this previously (Abraham, 1973–OM3; 1976–OM4) as have Komogorov *et al.* (1937), Rashevsky (1940), Turing (1952) Rosen (1970), Arbib (1972), Katchlasky and Neumann (1972), Hoffman (1977), Zeeman (1977: 293), Freeman (1981),

and many others. In this section, we review and expand the field scheme described earlier (Abraham, 1983*c*–OM11).

We think of a biological organ (for example, the hypothalamus) at once as a three-dimensional continuum of biophysical matter, and as an aggregation of cells. We carry along both of these images simultaneously, as in the wave/particle duality of physics. We suppose that each cell is reasonably well modeled by a simple dynamical scheme (dynamical system with controls). Admittedly, this is an extreme oversimplification. Beginning with an *excitation*, control metabolites (hormones, neuro-transmitters, morphogens, etc.) diffuse and react slowly in this continuum. On a shorter time scale, convection currents within the cells average the momentary concentrations of these control metabolites. On yet a shorter time scale, the cell dynamics move the physiological state of each cell through a brief transient, to the attractor determined by the average control metabolite concentrations, the states of neighboring cells, and its initial condition (previous attractor) at the time of excitation. [see Abraham and Shaw (1982) for an introduction to these concepts of dynamical systems theory, and 'Dynamical models for physiology' (Abraham, 1983*c*–OM11) for their application in this context.]

There may be different types of cells interspersed and matted in the aggregation. [see Hoffman (1977) for an interesting classification of those in the visual cortex, based on Lie algebra theory.] Even though this may be essential for our scheme, we will suppose now, to simplify the discussion, that they are all of one sort. (This might be justifiable in the case of slime mold, or the liver.) Later, in Section 10, we will relax this restriction. With this simplifying assumption, we may visualize the instantaneous state of the organ in our scheme as follows.

1. Choose a simple dynamical scheme modeling the standard cell of the organ, the *standard scheme*.

2. Discretize the domain, choosing a point *centrum* in each cell.

3. Represent the continuous distribution of control metabolites throughout the organ by a map from the physical domain into the control space of the standard scheme. This *continuous control field* may be regarded as a vector in an infinite-dimensional state space of the control system, in the language of complex dynamical systems.

4. Dually, represent the control metabolite levels at the centra by a map from the centra (actually, a finite set of indices identifying them) into the control space of the standard scheme. This *discrete control field* may be regarded as a finite-dimensional state of the reduced control system.

At this point we visualize the organ, its cellular decomposition, and its centra, all imaged in the control space of the standard scheme by these two maps. We proceed now by interpreting the dynamics of each cell in the standard scheme.

5. Choose an initial state for each cell, and visualize it in the standard scheme, hovering over the image in the control space of the centrum of that cell, under the control field. These choices, together, comprise the *initial field* of the organ, in this modeling scheme.

6. Start up the cellular dynamics, in a separate copy of the model for each cell, and wait for dynamic equilibrium in each.

As the transients die away, one attractor is occupied over each centrum. We obtain an image of the organ on the locus of attraction of the standard scheme, representing the *instantaneous state of the organ*. The initial field has evolved to the final attractor field. During this process, we have assumed that the control fields have not changed, or at least, that *they change very slowly* with respect to the internal dynamics of the standard cell.

To visualize the instantaneous state of the organ, we use the trick of Zeeman, and *observe the control fields and attractor fields in the standard scheme for a single cell*. If the control

metabolite levels in the organ slowly change, the control and attractor fields move about in the standard scheme. Thus, the attractor field will shlep along the locus of attraction. Whenever the control field transits the bifurcation set in the control space of the standard scheme, the attractor field will transit a bifurcation (subtle or catastrophic) of attractors, in the locus of attraction.

In summary, under all these assumptions, the current state of the organ is represented by the attractor field. This is a map from the cell centra to subsets (attractors) of the locus of attraction. This map *covers* the discrete control field. That is, to each centrum is assigned an attractor of the dynamical system determined by the control levels of that centrum.

In Fig. 2 one is drawn in three dimensions, assuming:

(a) the number of control metabolites is 2;

(b) the number of internal state variables of the standard cell is 1;

(c) the bifurcation diagram of the standard scheme is the cusp catastrophe, as used by Zeeman for his heart model;

(d) the dimension of the physical substrate is 2;

(e) the number of cells is 3.

The complex dynamical scheme

At this point we may draw the connection between this field approach, and the specifics of complex dynamical system theory.

The actual dynamical scheme for the finite set of cells is a Cartesian product of identical copies of the standard scheme, on copy for each cell. The dynamical evolution, and the single attractor representing the dynamical equilibrium of the combined system of all the cells, belongs to this, which we will call the *big scheme*. The discrete control field determines a single point in the control space of this combined system. The continuous control field is an instantaneous state of another dynamical system, of infinite dimension, which models the diffusion and reaction of control metabolites in the physical domain. This is the *executive*

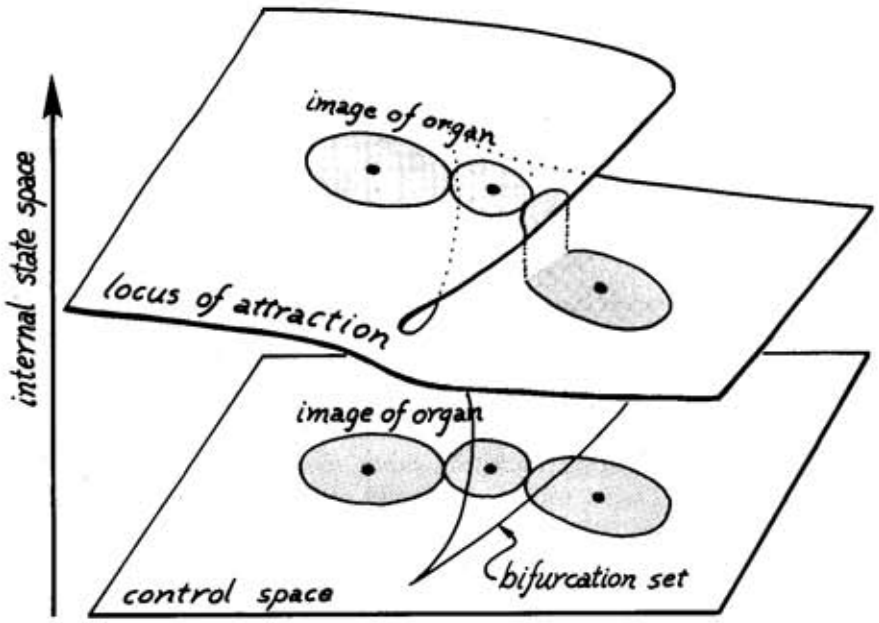

Fig. 2. Images of the organ in the dynamical scheme.

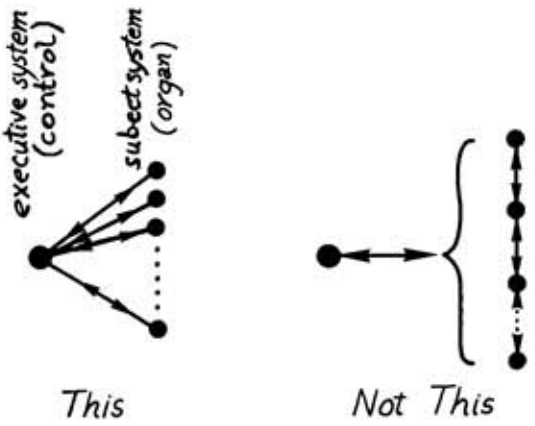

Fig. 3. Schematic of the serially bicoupled network.

scheme. The discrete control field represents the static scheme coupling the control metabolite diffusion system to the big scheme for the cellular organ, which is the *subject scheme*. The production of control metabolites by the cells represents a feedback from the subject scheme to the executive. Thus, we have a serially bicoupled network scheme for the total system, as shown, in Fig. 3.

The serial bicoupling between the executive scheme for the control metabolite system and the big scheme for the cellular organ replaces the usual idea of a direct parallel coupling between the cells. It can easily accommodate different types of coupling, such as: diffusion through tissue, blood circulation, pipette systems such as the portal-hypophyseal ventals, and synaptic junctions.

Our strategy in this paper is to fasten upon the attractor field, and forget about the actual attractor in the big scheme. This is reasonable, if the cells are not dynamically coupled to each other. This is not realistic (see Abraham, 1983*c*, sections B5, C4) but it will allow us to visualize the concepts under discussion, all of which apply equally to the more general scheme. A more realistic model for the organ would allow spatially modulated coupling between cells. Thus, the subject scheme would be a serially bicoupled network, as suggested by Kolmogorov *et al.* (see Abraham, 1983*c*). And now, as this model is too large to visualize, we return to the field approach.

The field scheme for thoughts

We suppose now that we have at hand a field model for the mammalian brain. An instantaneous state, a *thought*, is controlled by a control field, and is represented by an attractor field. Unlike our simplistic three-dimensional example in the preceding section, this one may have many control metabolite dimensions, and many internal state variables. Thus, the typical attractor will not be static (a point) or even periodic (a cycle) but most probably chaotic. Nevertheless, as multiple surfaces terminated by catastrophic bifurcations abound in typical bifurcation diagrams, hysteresis will be a principal feature of the standard scheme. Our emphasis here is on *the geometry of the locus of attraction*, not on the qualitative features of the individual attractors.

Thus, an attractor field is a thought. A thought covers a control field. Different thoughts can cover the same control field. Changing the control field will change the thoughts in a deterministic way. The (discrete) control field belongs to a finite-dimensional space, and its

change should be thought of as a curve in that space. But in the spirit of control theory, we will also think of such a curve of control fields, in discrete approximation, as a finite sequence of points. And beginning from the same initial control field, different sequences of intermediate control fields may end at the same final control field. And the same initial thought, shlepped along the locus of attraction by these different sequences of intermediate control fields, can end up as different thoughts, covering the same final control field. This is because there is *hysteresis* in the standard scheme. We think of the controlling sequence of control metabolite maps as a *program*, which means something like this in information theory. And we refer to this dependence of the change in the attractor field upon the program as *holonomy*, which means something like this in differential geometry.

Absolute programs

Note that the bifurcation set divides the control space (of virtual control metabolite concentrations) into a number of disjoint regions, *supposed finite*. Over each region there is a fixed set of attractors, also supposed finite. These are called *competing attractors* by Thom. Here, generic means that there is no centrum in the bifurcation set. Thus, for a fixed generic control field there is a finite number of attractor configurations possible, and in fact, a finite number for *any generic control field*. We regard them here as the *gamut of thoughts* possible over a generic control field. This is the basis for the following informatic metaphors.

An *address* is a generic control field. The *data* at that address is an attractor field (thought) covering it. A program, as defined above, is a sequence of addresses, regarded as the discrete approximation to a curve in the control space. We will call this an *absolute program*.

But to *copy data* to a final address, one must start with the right initial data at the initial address, and run the right program from the initial address to the final one. Changing either the initial data (thought, attractor field), or the program (sequence of addresses, control fields) ends at different data, although at the same address.

Integrative programs

But now to complete the connection between the neurophysiological model and the informatic metaphor, we must allow for the process of feedback from states (attractor fields) to controls (control fields), for *the cells may produce (or destroy) control metabolites*. Thus, we will generalize the idea of a program as follows. We introduce now a new hypothesis: *the control space of the standard cell is a vector space*. Thus control fields comprise a finite-dimensional vector space, and we may add them. This hypothesis can be easily generalized, by introducing a nonlinear map in place of addition, but this would only complicate the discussion unnecessarily.

Further, we introduce a new structure, to represent feedback from the subject system to the executive system. This is (in its simplest version) a function from the internal state space to the control space of the standard cell. We assume that in a given instantaneous state of the standard cell (represented by a point in the internal state space) control metabolites are produced at a constant rate. The new function, the *rate function*, specifies this rate.

Further, we assume that each cell (represented by a given concentration of control metabolites, and an attractor of the associated dynamical system) produces control metabolites at a constant rate, averaged over the fast variation of internal parameters along the attractor. Next, we suppose that sequential programs are run according to a strict clock. Thus, in each unit interval of time, corresponding to one instruction of the program, a cell in

a given attractor (data) over a given control metabolite level (address) will produce an increment of control metabolite which is the integral of the averaged rate function over the clock interval. Thus, given an instantaneous state of the entire organ, each cell produces increments of control metabolite. The contribution of each cell results in an increment to the control metabolite concentration for itself, and also its neighbors. So finally we must have in the scheme a spatial rule to specify how an increment of control metabolite concentration in the original organ will change the control field.

This should be expressed as a *dynamical system* on the finite-dimensional vector space of control fields, such as reaction-diffusion equations. In this case, we have this situation: there is a control field (original address) at an attractor of its own dynamic on the control space, a fast perturbation arrives, the control-space dynamic relaxes the perturbed state to the appropriate attractor (new address). If the perturbation does not push the instantaneous control state into a new basin, it relaxes to the orginal attractor: no address increment. In any case, we assume now such a rule: at each state of the organ, an increment to the control field is determined. We call this the *address increment* of the given state of the entire organ.

A program now will consist of initial data at an initial address, and a sequence of *relative addresses.* Each step of the program goes to a new address (control field on centra) determined by adding the current address, the address increment from the current state (assumed to be a constant) and the relative address at the current step of the program. Now, *the same program*, with different initial data, *can end at different addresses*. We will call this new kind of program an *integrative program.*

Holonomy programs

Finally, imagine a integrative program which begins and ends at the same *end field*. Its effect upon all the attractor fields (thoughts) covering its end field is to *map* them among themselves, generally in a many-to-one manner. This map is analogous to the holonomy concept of differential geometry. It is caused by the catastrophic bifurcations in the standard cell model.

We shall have use for some of the language of holonomy, from differential geometry. The invertible holonomy programs, at a given end field, comprise the *holonomy monoid* of that address. And for a given file (data, attractor field) at that address (end field), the set of all files obtainable from it by the operation of invertible holonomies upon it comprise the *holonomy orbit* of the original file.

The holonomy of bifurcation diagrams could (and perhaps will) be studied in the abstract. Meanwhile, it will be fundamental in our application of the field scheme for the brain to cognitive processes, in the next section.

Abstraction and application

Now we are ready to apply our scheme to the abstraction and application processes described in a preceding section, in the informatic metaphor. The main idea is that *the abstract concept* on level M, in the hierarchical model of consciousness described previously, *is an initial data file* (thought, initial attractor field) at a given address, as described in the preceding section. This is a filename, or pointer program. That is, it is a dynamical state generating an integrative program leading to a key state on level S. Its *instances* (thoughts, as final attractor fields) on level S, *are the files of the holonomy orbit* of this *key state*. The *name* of an instance on level S is the holonomy which creates it from the key state, which represents the abstraction on level S.

The *old application process* just requires running the appropriate holonomy program, its name, starting from the correct initial address and data of the abstraction. The *new application process* is a little more difficult. It requires the recognition of an existing abstraction, of which the new file is an instance. That is, an initial file (key of the abstraction) must be found, from among those already learned and ensconced on level M, and a holonomy program from key to instance. Thus, recognition is carried out, in this scheme, by running trial holonomies, hoping to strike the key of an existing abstract model of level M. If found, the recognition problem is solved. (The difficulty in finding one is resolved in the next section). The successful holonomy program must now be *inverted*, to name the new instance.

Two problems are encountered here:

1. How to determine, within the neurophysiological model, the close approach of two files. That is, while moving a file by a holonomy program, when is it close to another file (for example, an existing abstract model), and converging to it. We call this the *convergence problem*.
2. Having found an interesting holonomy program, how can we find an inverse holonomy, if there is one. We call this the *inversion problem*.

The *abstraction* process is even more difficult. From several instances on level S, one seeks integrative programs with identical final addresses, uses these programs to copy all the instances to this new address (this moves the file folders into a common pile), and operates on these transposed instances with the holonomy monoid of the common address. If possible, an abstraction (file on level M) will be found after enough trials, in the orbit of which lie all of the several instances. This presents, again, the two problems of *inversion* and *convergence*. In the next section, we describe a transformation from this scheme to another (the local geometrical model) in which the convergence problem is solved.

The geometric model

What we wish to do now is to simplify programs as much as possible, by standardizing the addresses. We observe that, although there is a continuum of addresses (vector space of control fields), most of them are equivalent. That is, any address can be deformed into another without essentially changing the data (attractor field) as long as *no centrum crosses a hypersurface* of the bifurcation set. We chose, thus, a *standard control point* in each component of the regular set (complement of the bifurcation set) in the control space. At the end of each program, we deform the final address into a standard one, by standardizing the control metabolite concentration at each centrum. Thus if a centrum has control metabolite concentration in one of the components, we deform this concentration to the standard one, *within the same* component, as shown in Fig. 4. For the moment, we regard this as a mathematical transformation only. Later, we will propose a neurophysiological mechanism for this, inspired by Hoffman.

A *standard field* is identified by two labels attached to each cell: the component of control space occupied by its (standardized) control metabolite concentration, and its attractor. These comprise a *standard state* for that cell. What we have achieved, through standardization, is a physical location, in the state space of the cell, for each of the possible attractors of a given regular component of control space. In fact, choosing one point in each basin (we think of this as the *average* state of the cell, for the given attractor) we may define the *distance* between two standard fields, as shown in Fig. 5. And thus, we can try to minimize this distance. For standard configurations with separate addresses, we say the distance is infinite.

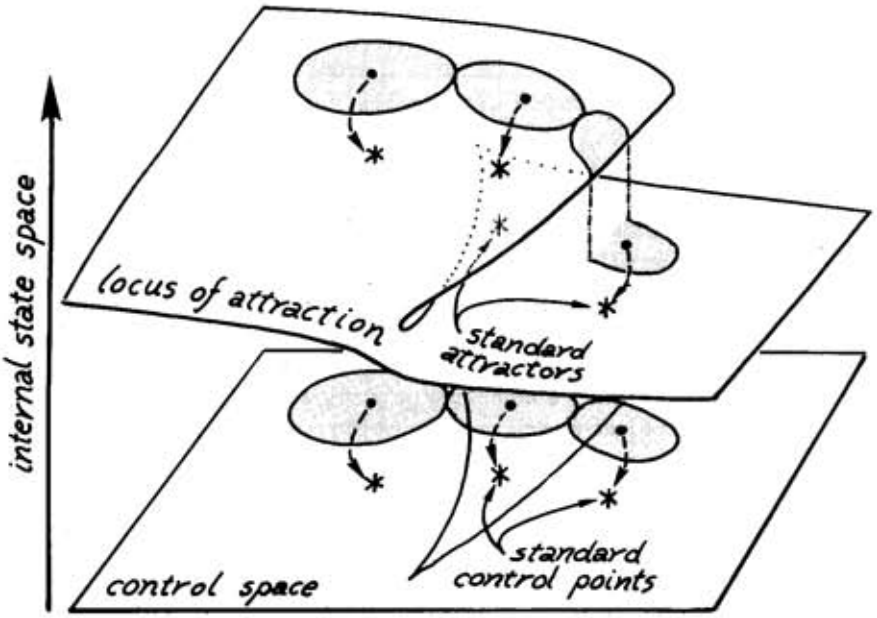

Fig. 4. Deformation to a standard field.

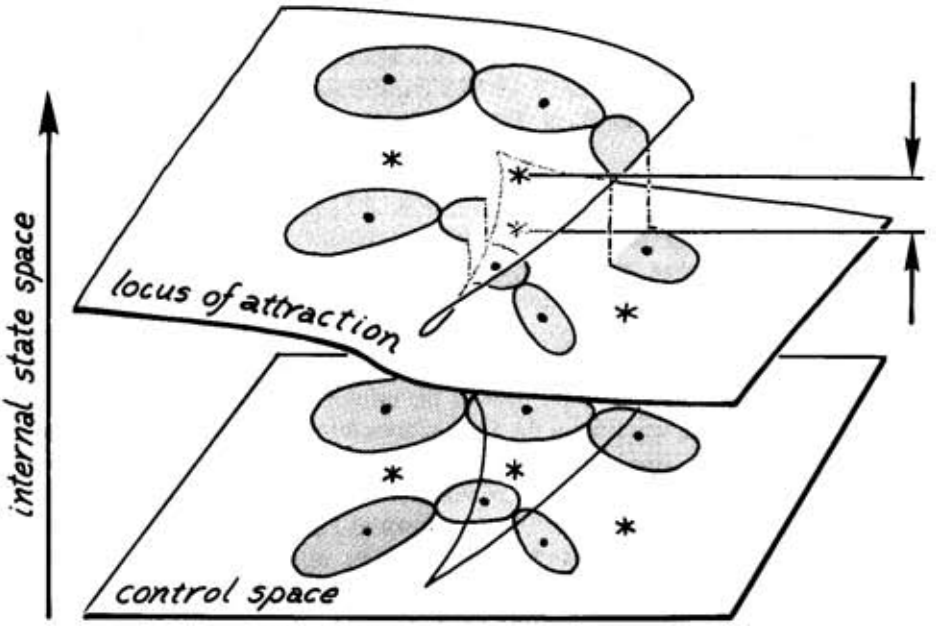

Fig. 5. Distance between two fields.

Now, at last, we can imagine a possible algorithm for the abstraction process. Given several *instances*, we standardize them. In the simpler cases, they comprise different data over the same address. Otherwise, we copy the data to the same address with an integrative program. If there is no intergrative program from an instance to a common address, it must be abandoned.

Now, as described in the preceding section, we experiment with our favorite *holonomies*, integrative programs which begin and end at the same address. From all the different initial

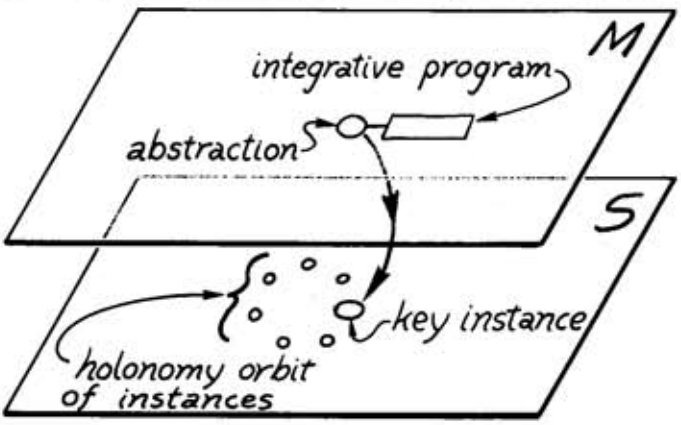

Fig. 6. Abstraction, key, and instance orbit.

data, we are seeking common final data. So for each instance file, we run all of our holonomies, generating its *instance orbit*. Files (attractor fields) in the intersection of these instance orbits are candidate abstraction keys for these instances. They must be tried, one at a time.

For each candidate, we apply the holonomy monoid to create its orbit. Then, the distances from each instance to the orbit must be added, to measure the value of the candidate abstraction. Varying the candidate, we seek to minimize this distance. If successful, we have found an abstraction for the given instances. The abstraction on level M is an integrative program leading to the key.

We think of the instances, standardized to the same address, as the pile of file folders, in the earlier mechanical metaphor. But in this case, it is not the address which is the abstraction, but the file on level M, which points to a key file in the pile, as shown in Fig. 6.

Special purpose cells

Finally, we will describe the Hoffman-inspired neurophysiological mechanism for the standardization of addresses. This requires adding some auxiliary cells to the neurophysiological scheme. Thus we will have a complex dynamical scheme for the brain which has a simple standard cell, inhomogeneous cell-types, and nonuniform control metabolite distribution, in the framework of organic resolution (Abraham 1983*c*–OM11).

First, there must be *buffer cells*. Their purpose is to sense nearby control metabolite concentrations, and buffer them, simultaneously inhibiting any efforts of uniform cells to change the address (control metabolite concentration) as if in response to excitation. They have to steer away from the bifurcation set.

We may imagine a gradient-like dynamical system in the control space, moving away from the bifurcation set toward a distinguished central point of each component (see Fig. 7). But recall that an address is a control field, and the effect of buffer cells is to pull the image of the (cellular) physical substrate into these sinks. The image is stretched tight across the bifurcation set (see Fig. 8). The high gradients between cells on these boundaries will overpower buffering, and limit gradients will exist. A typical control field is shown as a graph in Fig. 9. So we propose an additional distribution of special purpose cells, the sample-and-hold cells. These sense a region of essentially constant control metabolite levels, sample these levels, and hold them. The geometric evaluation of candidate abstraction programs, described above, is accomplished by the sample-and-hold cells.

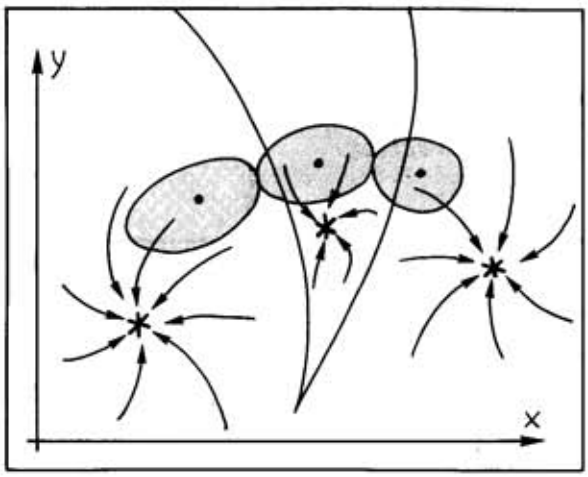

Fig. 7. Standardization of control field, in progress.

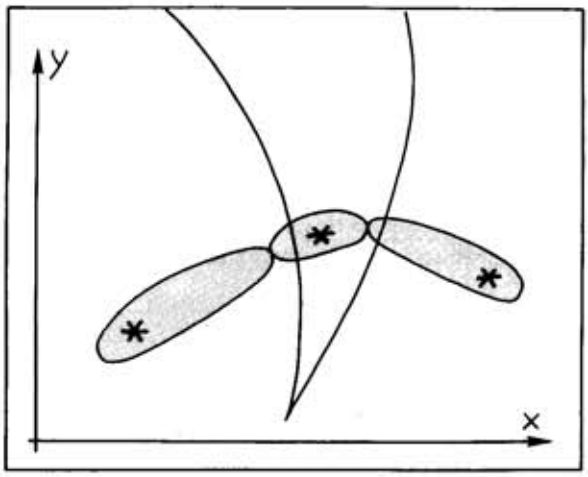

Fig. 8. Standardization of control field, completed.

Conclusion

Here we have described a hypothetical scheme for communication between adjacent levels in an hierarchical information structure of a conscious mind. In informatic and mechanical metaphors, we have described the processes:

(a) *old application* of an existing abstraction to an existing instance:
(b) *new application* of an existing abstraction to a new instance: and
(c) *new abstraction* of similar instances into a new abstraction.

The informatic metaphors developed here are:

(1) *address* as a standard control field;
(2) *data* as an attractor field;
(3) *data comparison* by a geometric distance;

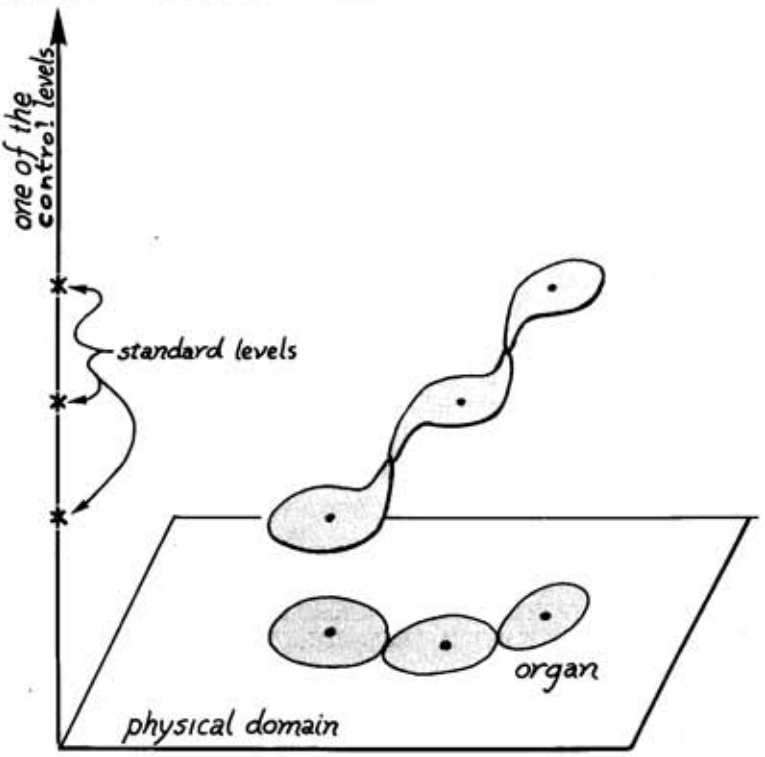

Fig. 9. Graph of standard control field.

(4) *program* as a sequence of relative address moves, integrating data at each incremental address, according to a dynamical system on the address space, and dragging the data along the locus of attraction of the dynamical scheme for the standard cell, and
(5) *filename* (*pointer*) expansion by a integrative program.

Further, an homogeneous neurophysiological model has been proposed for the realization of these processes and metaphors in mammalian brains. The application process is simple and direct in this scheme, based on the informatic metaphor. The abstraction process is more complicated. For this, we have introduced two supplementary distributions of special cells: buffer cells and sample-and-hold cells. In this inhomogeneous neurophysiological model, the mechanical metaphor (aggregation of files to a pile of a common address) becomes simple and direct as well.

All this is based on the theory of complex dynamical systems, developed in the earlier papers of this series (OM8–11: Abraham & Shaw, 1983; Abraham, 1983*a,b,c*) and developed in the context of the bifurcation diagram of an *imaginary neuron model*. This theory provides some guidance, even when the actual dynamical model is unknown. For this reason, we call it a *scheme*, rather than a model, for thought.

Nevertheless, one could go much further with this theory if explicit dynamical models (for a single neuron, for example) were known. In addition, the existence of buffer and sample-and-hold cells is hypothetical here. This part of the theory is inspired by the work of Hoffman, which goes much further is studying the function of special distributions of cells. We should like a more concrete neurophysiological proposal for these cells. The structure of the glial body, as a bundle over the cortex, is suggestive here. We propose these problems to neurobiologists.

Acknowledgements

This paper, written on 17 October, 1981 on the train from Paris to Amsterdam in a moment of revelation, is the translation into the new language of complex dynamics of some old ideas, familiar to visionaries of the past and present. In particular, Arbib, Freeman, Hoffman, Katchalsky, Kolmogorov, Plato, Rashevsky, Rosen, Thom, Turing, and Zeeman are mentioned in the text. Here I would like to recognize the particularly seminal work of Aharon Katzir-Katchalsky, which appeared shortly after his death, at the beginning of my interest in this subject. Finally, I would like to acknowledge the influence of important conversations with Fred Abraham, Walter Freeman, Alan Garfinkel, William Hoffman, Tim Poston, Dan Sunday, Rene Thom, and Christopher Zeeman, the support of the Institut des Hautes Etudes Scientifiques, particularly Nicolaas Kuiper, during the preparation of this work, and valuable feedback from the referee on expository matters.

References

Abraham, R. H. (1973). In: *Proceedings, First International Congress of Psychotronics and Paraphsychology*. Prague. (OM1).

Abraham, R. H. (1976). In (E. Jantsch & C. H. Waddington, Eds) *Evolution and Consciousness*. Reading, MA: Addison Wesley (OM4).

Abraham, R. H. (1976). In (P. Hilton, Ed.) *Structural Stability, the Theory of Catastrophes, and Applications*, Lecture Notes in Mathematics, 525. New York: Springer-Verlag. (OM5).

Abraham, R. H. (1981). In (E. Jantsch. Ed.) *The Evolutionary Vision*, AAAS Selected Symposium, 61. (OM7).

Abraham, R. H. (1983*a*). In (F. E. Yates, Ed.) *Self-organizing systems*. New York: Plenum (to appear) (OM9).

Abraham, R. H. (1983*b*). In (T. M. Rassias, Ed.) *Global Analysis-Analysis on Manifolds* Leibzig: (to appear) (OM10).

Abraham, R. H., (1983*c*). *Am. J. Physiol.* (to appear). (OM11).

Abraham, R. H. & Shaw, C. D. (1982). *Dynamics, the Geometry of Behavior, Part 1: Periodic Behavior*. Santa Curz, CA: Aerial.

Abraham, R. H. & Shaw, C. D. (1983). In: F. E. Yates (Ed.) *Self-Organizing Systems*. New York: Plenum. (to appear). (OM8).

Arbib, M. (1972). *The Metaphorical Brain* New York: Interscience

Freeman, W. J. (1981) *Perspect Biol. Med.* (summer).

Hoffman, W. C. (1977). *Cahiers de Psychologie* **29**, 135–174.

Katchalsky, A. & Neumann E. (1972). *Int. J. Neurosci.* **3**, 175–182.

Kolmogorov, A., Petrovski I. & Pikunov, N. (1937). Etude de l'equation de la diffusion avec croissance de la quantite de matiere at son application a une probleme biologique, *Bull. Univ. d'Etat Moscou, Ser. Intern.*, **1** (A), 1–25.

Rashevsky (1940).

Rosen, R. (1970), *Dynamical System Theory in Biology*, Vol. 1. New York: John Wiley.

Shear, J. (1978), Plato, Piaget, and Maharishi on cognitive development (preprint).

Smith, W. R. (1980). *Bull. Math. Biol.* **42**, 57–78.

Thom, R. (1972). *Structural Stability and Morphogenesis, an Outline of a General Theory of Models*. English translation 1975 by D. Fowler. Reading, MA: Benjamin-Cummings.

Turing, (1952).

Zeeman, E. C. (1977). Catastrophe Theory and its Applications. Reading, MA: Addison-Wesley.